

# UCB Algorithm for Exponential Distributions

Wassim Jouini and Christophe Moy

## Abstract

We introduce in this paper a new algorithm for Multi-Armed Bandit (MAB) problems. A machine learning paradigm popular within Cognitive Network related topics (e.g., Spectrum Sensing and Allocation). We focus on the case where the rewards are exponentially distributed, which is common when dealing with Rayleigh fading channels. This strategy, named Multiplicative Upper Confidence Bound (MUCB), associates a utility index to every available arm, and then selects the arm with the highest index. For every arm, the associated index is equal to the product of a multiplicative factor by the sample mean of the rewards collected by this arm. We show that the MUCB policy has a low complexity and is order optimal.

## Index Terms

Learning, Multi-armed bandit, Upper Confidence Bound Algorithm, UCB, MUCB, exponential distribution.

## I. INTRODUCTION

Several sequential decision making problems face a dilemma between the exploration of a space of choices, or solutions, and the exploitation of the information available to the decision maker. The problem described herein is known as sequential decision making under uncertainty. In this paper we focus on a sub-class of this problem, where the decision maker has a discrete set of stateless choices and the added information is a real valued sequence (of feedbacks, or rewards) that quantifies how well the decision maker behaved in the previous time steps. This particular instance of sequential decision making problems is generally known as the multi-armed bandit (MAB) problem [1], [2].

A common approach to solving the exploration versus exploitation dilemma within MAB problems consists in assigning an utility value to every arm. An arm's utility aggregates all the past information

SUPELEC, IETR, SCEE, Avenue de la Boulaie, CS 47601, 35576 Cesson Sévigné, France.

Email: wassim.jouini@supelec.fr

about the lever and quantifies the gambler's interest in pulling it. Such utilities are called *indexes*. Agrawal et al. [2] emphasized the family of indexes minimizing the expected cumulated loss and called them Upper Confidence Bound (UCB) indexes. UCB indexes provide an optimistic estimation of the arms' performances while ensuring a rapidly decreasing probability of selecting a suboptimal arm. The decision maker builds its policy by greedily selecting the largest index. Recently, Auer et al. [3] proved that a simple additive form, of the rewards' sample mean and a bias, known as  $UCB_1$  can achieve order optimality over time when dealing with rewards drawn from bounded distributions. Tackling exponentially distributed rewards remains however a challenge as optimal learning algorithms to tackle this matter prove to be complex to implement [1], [2].

This paper is inspired from the aforementioned work. However, we suggest the analysis of a multiplicative rather than an additive expression for the index.

The main contribution of this paper is to design and analyze a simple, deterministic, multiplicative index-based policy. The decision making strategy computes an index associated to every available arm, and then selects the arm with the highest index. Every index associated to an arm is equal to the product of the sample mean of the reward collected by this arm and a scaling factor. The scaling factor is chosen so as to provide an optimistic estimation of the considered arm's performance.

We show that our decision policy has both a low computational complexity and can lead to a logarithmic loss over time under some non-restrictive conditions. For the rest of this paper we will refer to our suggested policy as *Multiplicative Upper Confidence Bound index* (MUCB).

The outline of this paper is the following: We start by presenting some general notions on the multi-armed bandit framework with exponentially distributed rewards in Section II. Then, Section III introduces our index policy and Section IV analyzes its behavior, proving the order optimality of the suggested algorithm. Finally, Section V concludes.

## II. MULTI-ARMED BANDITS

A  $K$ -armed bandit ( $K \in \mathbb{N}$ ) is a machine learning problem based on an analogy with the traditional slot machine (one-armed bandit) but with more than one lever. Such a problem is defined by the  $K$ -tuple  $(\theta_1, \theta_2, \dots, \theta_K) \in \Theta^K$ ,  $\Theta$  being the set of all positive reward distributions. When pulled at a time  $t \in \mathbb{N}$ , each lever<sup>1</sup>  $k \in \llbracket 1, K \rrbracket$  (where  $\llbracket 1, K \rrbracket = \{1, \dots, K\}$ ) provides a reward  $r_t$  drawn from a distribution  $\theta_k$  associated to that specific lever. The objective of the gambler is to maximize the cumulated sum of rewards

<sup>1</sup>We use indifferently the words “lever”, “arm”, or “machine”.

through iterative pulls. It is generally assumed that the gambler has no (or partial) initial knowledge about the levers. The crucial tradeoff the gambler faces at each trial is between *exploitation* of the lever that has the highest expected payoff and *exploration* to get more information about the expected payoffs of the other levers. In this paper, we assume that the different exponentially distributed payoffs drawn from a machine are independent and identically distributed (i.i.d.) and that the independence of the rewards holds between the machines. However the different machines' reward distributions  $(\theta_1, \theta_2, \dots, \theta_K)$  are not supposed to be the same.

Let  $I_t \in \llbracket 1, K \rrbracket$  denote the machine selected at a time  $t$ , and let  $H_t$  be the history vector available to the gambler at instant  $t$ , i.e.,  $H_t = [I_0, r_0, I_1, r_1, \dots, I_{t-1}, r_{t-1}]$

We assume that the gambler uses a policy  $\pi$  to select arm  $I_t$  at instant  $t$ , such that  $I_t = \pi(H_t)$ . We shall also write  $\forall k \in \llbracket 1, K \rrbracket$ ,  $\mu_k \triangleq \frac{1}{\lambda_k} \triangleq \mathbb{E}[\theta_k]$ , where  $\lambda_k$  refers to the parameter of the considered exponential distribution with pdf  $f_{\theta_k}(x) = \lambda_k e^{-\lambda_k x}$ ,  $x \geq 0$ , and we assume that  $\mu_k > 0, \forall k \in \llbracket 1, K \rrbracket$ . The (cumulated) regret of a policy  $\pi$  at time  $t$  (after  $t$  pulls) is defined as follows:  $R_t = t\mu^* - \sum_{m=0}^{t-1} r_m$ , where  $\mu^* = \max_{k \in \llbracket 1, K \rrbracket} \{\mu_k\}$  refers to the expected reward of the optimal arm.

We seek to find a policy that minimizes the expected cumulated regret (Equation 1),

$$\mathbb{E}[R_t] = \sum_{k \neq k^*} \Delta_k \mathbb{E}[T_{k,t}], \quad (1)$$

where  $\Delta_k = \mu^* - \mu_k$  is the expected loss of playing arm  $k$ , and  $T_{k,t}$  refers to the number of times the machine  $k$  has been played from instant 0 to instant  $t - 1$ .

### III. MULTIPLICATIVE UPPER CONFIDENCE BOUND ALGORITHMS

This section presents our main contribution, the introduction of a new multiplicative index. Let  $B_{k,t}(T_{k,t})$  denote the index of arm  $k$  at time  $t$  after being pulled  $T_{k,t}$ . We refer to as Multiplicative Upper Confidence Bound algorithms (MUCB) the family of indexes that can be written in the form:

$$B_{k,t}(T_{k,t}) = \overline{X}_{k,t}(T_{k,t}) M_{k,t}(T_{k,t}),$$

where  $\overline{X}_{k,t}(T_{k,t})$  is the sample mean of machine  $k$  at step  $t$  after  $T_{k,t}$  pulls, i.e.,  $\overline{X}_{k,t}(T_{k,t}) = \frac{1}{T_{k,t}} \sum_{i=0}^{t-1} \mathbb{1}_{\{I_i=k\}} r_i$  and  $M_{k,t}(\cdot)$  is an upper confidence scaling factor chosen to insure that the index  $B_{k,t}(T_{k,t})$  is an increasing function of the number of rounds  $t$ . This last property insures that the index of an arm that has not been pulled for a long time will increase, thus eventually leading to the sampling of this arm. We introduce a

particular parametric class of MUCB indexes, which we call  $MUCB(\alpha)$ , given as follows<sup>2</sup>:

$$\forall \alpha \geq 0, \quad M_{k,t}(T_{k,t}) = \frac{1}{\max \left\{ 0; \left( 1 - \sqrt{\frac{\alpha \ln(t)}{T_{k,t}}} \right) \right\}} \quad (2)$$

We adopt the convention that  $\frac{1}{0} = +\infty$ . Given a history  $H_t$ , one can compute the values of  $T_{k,t}$  and  $M_{k,t}$  and derive an index-based policy  $\pi$  as follows:

$$I_t = \pi(H_t) \in \arg \max_{k \in \llbracket 1, K \rrbracket} \{B_{k,t}(T_{k,t})\}. \quad (3)$$

#### IV. ANALYSIS OF $MUCB(\alpha)$ POLICIES

This section analyses the theoretical properties of  $MUCB(\alpha)$  algorithms. More specifically, it focuses on determining how fast is the optimal arm identified and what are the probabilities of *anomalies*, that is sub-optimal pulls.

##### A. Consistency and order optimality of MUCB indexes

*Definition 1 ( $\beta$ -consistency):* Consider the set  $\Theta^K$  of  $K$ -armed bandit problems. A policy  $\pi$  is said to be  $\beta$ -consistent,  $0 < \beta \leq 1$ , with respect to  $\Theta^K$ , if and only if

$$\forall (\theta_1, \dots, \theta_K) \in \Theta^K, \quad \lim_{t \rightarrow \infty} \frac{\mathbb{E}[R_t]}{t^\beta} = 0 \quad (4)$$

We expect good policies to be at least *1-consistent*. As a matter of fact, *1-consistency* ensures that, asymptotically, the average expected reward is optimal.

From the expression of Equation 1 one can remark that its is sufficient to upper bound the expected number of times  $\mathbb{E}[T_{k,t}]$  one plays a suboptimal machine  $k$  after  $t$  rounds, to obtain an upper bound on the expected cumulated regret. This leads to the following theorem.

*Theorem 1 (Order optimality of  $MUCB(\alpha)$  policies):* Let  $\rho_k = \mu_k/\mu^*$ ,  $k \in \llbracket 1, K \rrbracket \setminus \{k^*\}$ . For all  $K \geq 2$ , if policy  $MUCB(\alpha > 4)$  is run on  $K$  machines having rewards drawn from exponential distributions  $\theta_1, \dots, \theta_K$  then:

$$\mathbb{E}[R_t] \leq \sum_{k: \Delta_k > 0} \frac{4\mu^*\alpha}{1 - \rho_k} \ln(t) + o(\ln(t)) \quad (5)$$

<sup>2</sup>This form offers a compact mathematical formula. However practically speaking, a machine  $k$  is played when  $T_{k,t} \leq \alpha \ln(t)$ . Otherwise the machine with largest finite index is played.

Proving Theorem 1 relies on three lemmas that we analyze and prove in the next subsection. The lemma 1 provides a general bound for the regret regardless of the policy considered. The expression is function of two probabilities related to learning anomalies. These anomalies depend on the learning algorithm. They are introduced and analyzed. Then through lemma 2 and 3 we upper bound them.

### B. Learning Anomalies and Consistency of MUCB policies

Let us introduce the set  $\mathbb{S} = \mathbb{N} \times \mathbb{R}$ ; then, one can write  $S_{k,t} = (T_{k,t}, B_{k,t}) \in \mathbb{S}$  the decision state of arm  $k$  at time  $t$ . We associate the product order to the set  $\mathbb{S}$ : for a pair of states  $S = (T, B) \in \mathbb{S}$  and  $S' = (T', B') \in \mathbb{S}$ , we write  $S \geq S'$  if and only if  $T \geq T'$  and  $B \geq B'$ .

*Definition 2 (Anomaly of type 1):* We assume that there exists at least one suboptimal machine, i.e.,  $\llbracket 1, K \rrbracket \setminus \{k^*\} \neq \emptyset$ . We call *anomaly of type 1*, denoted by  $\{\phi_1(u_k)\}_{k,t}^\pi$ , for a suboptimal machine  $k \in \llbracket 1, K \rrbracket \setminus \{k^*\}$ , and with parameter  $u_k \in \mathbb{N}$ , the following event:

$$\{\phi_1(u_k)\}_{k,t}^\pi = \{S_{k,t} \geq (u_k, \mu^*)\}.$$

*Definition 3 (Anomaly of type 2):* We refer to as *anomaly of type 2*, denoted by  $\{\phi_2\}_t^\pi$ , associated to the optimal machine  $k^*$ , the following event:

$$\{\phi_2\}_t^\pi = \{S_{k^*,t} < (\infty, \mu^*) \cap T_{k^*,t} \geq 1\}.$$

*Lemma 1 (Expected cumulated regret. Proof in VI-B):* Given a policy  $\pi$  and a MAB problem, let  $\mathbf{u} = [u_1, \dots, u_K]$  represent a set of integers, then the expected cumulated regret is upper bounded by:

$$\mathbb{E}[R_t] \leq \sum_{k \neq k^*} \Delta_k u_k + \sum_{k \neq k^*} \Delta_k \mathbb{P}_t(u_k)$$

with,  $\mathbb{P}_t(u_k) = \sum_{m=u_k+1}^t \left( \mathbb{P}(\{\phi_2\}_m^\pi) + \mathbb{P}(\{\phi_1(u_k)\}_{k,m}^\pi) \right)$

We consider the following values for the set  $\mathbf{u}$ , for all suboptimal arms  $k$ ,  $u_k(t) = \left\lceil \frac{4\alpha}{(1-\rho_k)^2} \ln(t) \right\rceil$ .

We show in the two following lemmas that for the defined set  $\mathbf{u}$  the anomalies are upper bounded by exponentially decreasing functions of the number of iterations.

*Lemma 2 (Upper bound of Anomaly 1. Proof in VI-C):* For all  $K \geq 2$ , if policy  $MUCB(\alpha)$  is run on  $K$  machines having rewards drawn from exponential distributions  $\theta_2, \dots, \theta_K$  then  $\forall k \in \llbracket 1, K \rrbracket \setminus \{k^*\}$ :

$$\mathbb{P}(\{\phi_1(u_k)\}_{k,t}^\pi) \leq t^{-\alpha/2+1} \quad (6)$$

*Lemma 3 (Upper bound of Anomaly 2. Proof in VI-D):* For all  $K \geq 2$ , if policy  $MUCB(\alpha)$  is run on  $K$  machines having rewards drawn from exponential distributions  $\theta_1, \dots, \theta_K$  then:

$$\mathbb{P}(\{\phi_2\}_t^\pi) \leq t^{-\alpha/2+1} \quad (7)$$

We end this paper by the proof of Theorem 1.

*Proof of Theorem 1:* For  $\alpha > 4$ , relying on Lemmas 1, 2 and 3 we can write:

$$\mathbb{E}[R_t] \leq \sum_{k \neq k^*} \Delta_k \left\lceil \frac{4\alpha}{(1 - \rho_k)^2} \ln(t) \right\rceil + o(\ln(t))$$

with,  $\sum_{k \neq k^*} \Delta_k \mathbb{P}_t(u_k) = o(\ln(t))$ . Finally, since  $\Delta_k = \mu^*(1 - \rho_k)$  and  $u_k(t) = \frac{4\alpha}{(1 - \rho_k)^2} \ln(t) + o(\ln(t))$ , we find the stated result in Theorem 1. ■

## V. CONCLUSION

A new low complexity algorithm for MAB problems is suggested and analyzed in this paper: MUCB. The analysis of its regret proves that the algorithm is order optimality over time. In order to quantify its performance compared to optimal algorithms, further empirical evaluations are needed and are currently under investigation.

## ACKNOWLEDGMENT

The authors would like to thank Damien Ernst, Raphael Fonteneau and Emmanuel Rachelson for their many helpful comments and answers regarding this work.

## REFERENCES

- [1] T.L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6:4–22, 1985.
- [2] R. Agrawal. Sample mean based index policies with  $O(\log(n))$  regret for the multi-armed bandit problem. *Advances in Applied Probability*, 27:1054–1078, 1995.
- [3] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite time analysis of multi-armed bandit problems. *Machine learning*, 47(2/3):235–256, 2002.
- [4] H. Chernoff. A measure of asymptotic efficiency fo tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, pages 493–507, 1952.

## VI. APPENDIX

### A. Large deviations inequalities

*Assumption 1 (Cramer condition):* Let  $X$  be a real random variable.  $X$  satisfies the Cramer condition if and only if

$$\exists \gamma > 0 : \forall \eta \in (0, \gamma), \mathbb{E} [e^{\eta X}] < \infty .$$

*Lemma 4 (Cramer-Chernoff Lemma for the sample mean):* Let  $X_1, \dots, X_n$  ( $n \in \mathbb{N}$ ) be a sequence of i.i.d. real random variables satisfying the Cramer condition with expected value  $\mathbb{E}[X]$ . We denote by  $\bar{X}_n$  the sample mean  $\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$ . Then, there exist two functions  $l_1(\cdot)$  and  $l_2(\cdot)$  such that:

$$\forall \beta_1 > \mathbb{E}[X], \mathbb{P}(\bar{X}_n \geq \beta_1) \leq e^{-l_1(\beta_1)n} ,$$

$$\forall \beta_2 < \mathbb{E}[X], \mathbb{P}(\bar{X}_n \leq \beta_2) \leq e^{-l_2(\beta_2)n} .$$

Functions  $l_1(\cdot)$  and  $l_2(\cdot)$  do not depend on the sample size  $n$  and are continuous non-negative, strictly increasing (respectively strictly-decreasing) for all  $\beta_1 > \mathbb{E}(X)$  (respectively  $\beta_2 < \mathbb{E}(X)$ ), both null for  $\beta_1 = \beta_2 = \mathbb{E}(X)$ .

This result was initially proposed and proved in [4]. The bounds provided by this lemma are called *Large Deviations Inequalities* (LDIs) in this paper.

In the case of exponential distributions this theorem can be applied and LDI functions have the following expressions:

$$l_1(\beta) = l_2(\beta) = \frac{\beta}{\mathbb{E}[X]} - 1 - \ln \left( \frac{\beta}{\mathbb{E}[X]} \right) \geq \frac{3 \left( 1 - \frac{\beta}{\mathbb{E}[X]} \right)^2}{2 \left( 1 + 2 \frac{\beta}{\mathbb{E}[X]} \right)}$$

### B. Proof of Lemma 1

According to Equation 1:  $\mathbb{E}[R_t^\pi] = \sum_{k \neq k^*} \Delta_k \mathbb{E}[T_{k,t}]$ . Per definition  $T_{k,t} = \sum_{m=0}^{t-1} \mathbf{1}_{I_m=k}$ . Then,  $\mathbb{E}[T_{k,t}] = \sum_{m=0}^{t-1} \mathbb{E}[\mathbf{1}_{I_m=k}]$ . After playing an arm  $u_k$  times, bounding the first  $u_k$  terms by 1 yields:

$$\mathbb{E}[T_{k,t}] \leq u_k + \sum_{m=u_k+1}^{t-1} \mathbb{P}(\{I_m = k\} \cap \{T_{k,m} > u_k\}) \quad (8)$$

Then we can notice that the following events are equivalent:

$$\{I_m = k\} = \left\{ B_{k,m} > \max_{k' \neq k} B_{k',m} \right\}$$

Moreover we can notice that:

$$\left\{ B_{k,m} > \max_{k' \neq k} B_{k',m} \right\} \subset \{ B_{k,m} > B_{k^*,m} \}$$

Which can be further included in the following union of events:

$$\{ B_{k,m} > B_{k^*,m} \} \subset \{ B_{k,m} > \mu^* \} \cup \{ \mu^* > B_{k^*,m} \}$$

Consequently we can write:

$$\{ I_m = k \} \cap \{ T_{k,m} > u_k \} \subset \{ \Phi_1(u_k) \}_{k,m}^\pi \cup \{ \Phi_2 \}_m^\pi \quad (9)$$

Finally, we apply the probability operator:

$$\mathbb{E}[T_{k,t}] \leq u_k + \sum_{m=u_k+1}^{t-1} \mathbb{P}(\{ \Phi_1(u_k) \}_{k,m}^\pi) + \mathbb{P}(\{ \Phi_2 \}_m^\pi). \quad (10)$$

The combination of Equation 1 - given at the beginning of this proof - and Equation 10 concludes this proof.

### C. Proof of Lemma 2

From the definition of  $\{ \phi_1(u_k) \}_{k,t}^\pi$  we can write that:

$$\begin{aligned} \mathbb{P}(\{ \phi_1(u_k) \}_{k,t}^\pi) &= \sum_{S_{k,t} \in \mathbb{S}} \mathbb{P}(S_{k,t} \geq (u_k, \mu^*)), \\ &\leq \sum_{u=u_k}^{t-1} \mathbb{P}(B_{k,t}(u) \geq \mu^*). \end{aligned}$$

In the case of MUCB policies, we have:

$$\forall u \leq t, \mathbb{P}(B_{k,t}(u) \geq \mu^*) = \mathbb{P}\left(\overline{X}_{k,t}(u) \geq \frac{\mu^*}{M_{k,t}(u)}\right)$$

Consequently, we can upper bound the probability of occurrence of type 1 anomalies by:

$$\mathbb{P}(\{ \phi_1(u_k) \}_{k,t}^\pi) \leq \sum_{u=u_k}^{t-1} \mathbb{P}\left(\overline{X}_{k,t}(u) \geq \frac{\mu^*}{M_{k,t}(u)}\right).$$

Let us define  $\beta_{k,t}(T_{k,t}) = \frac{\mu^*}{M_{k,t}(T_{k,t})}$ .

Since we are dealing with exponential distributions, the rewards provided by the arm  $k$  satisfy the Cramer condition. As a matter of fact, since  $u \geq u_k \geq \alpha \frac{\ln(t)}{(1-\rho_k)^2}$  then:

$$\beta_{k,t}(u) \lambda_k = \rho_k^{-1} \left( 1 - \sqrt{\alpha \frac{\ln(t)}{u}} \right) \geq 1$$



So, according to the large deviation inequality for  $\overline{X}_{k,t}(T_{k,t})$  given by Lemma 4 (with  $T_{k,t} \geq u_k$  and  $u_k$  large enough), there exists a continuous, non-decreasing, non-negative function  $l_{1,k}$  such that:

$$\mathbb{P}(\overline{X}_{k,t}(T_{k,t}) \geq \beta_{k,t}(T_{k,t}) | T_{k,t} = u) \leq e^{-l_{1,k}(\beta_{k,t}(u))u}.$$

Finally:

$$\mathbb{P}(\{\phi_1(u_k)\}_{k,t}^\pi) \leq \sum_{u=u_k}^{t-1} e^{-l_{1,k}(\beta_{k,t}(u))u}. \quad (11)$$

The end of this proof aims at proving that for  $u \geq u_k$ :  $l_{1,k}(\beta_{k,t}(u)) \geq \alpha \frac{\ln(t)}{2u}$ .

Note that since we are dealing with exponential distributions we can write:  $l_{1,k}(\beta_{k,t}(u)) \geq \frac{3(1-\beta_{k,t}(u)\lambda_k)^2}{2(1+2\beta_{k,t}(u)\lambda_k)}$ .

Moreover since  $u \geq u_k \geq \alpha \frac{\ln(t)}{(1-\rho_k)^2}$  then:

$$\beta_{k,t}(u)\lambda_k = \rho_k^{-1} \left( 1 - \sqrt{\alpha \frac{\ln(t)}{u}} \right) \leq \rho_k^{-1}$$

Consequently it is sufficient to prove that:

$$\frac{3(1-\beta_{k,t}(u)\lambda_k)^2}{2(1+2\rho_k^{-1})} \geq \alpha \frac{\ln(t)}{2u}$$

Let us define  $h(t)$  as a function of time:  $h(t) = \sqrt{\alpha \frac{\ln(t)}{u}} \in [0, 1]$ . We analyze the sign of the function:

$$g(t) = (\rho_k^{-1}h(t) - (\rho_k^{-1} - 1))^2 - \frac{(1+2\rho_k^{-1})}{3}h(t)^2 \quad (12)$$

Consequently we need to prove that for  $u \geq u_k$ ,  $g(\cdot)$  has positive values.

Factorizing last equation leads to the following to terms:

$$\begin{cases} \left( \rho_k^{-1} - \sqrt{\frac{(1+2\rho_k^{-1})}{3}} \right) h(t) - (\rho_k^{-1} - 1) \\ \left( \rho_k^{-1} + \sqrt{\frac{(1+2\rho_k^{-1})}{3}} \right) h(t) - (\rho_k^{-1} - 1) \end{cases} \quad (13)$$

Since per definition:  $h(t) \in [0, 1]$  and  $\rho_k^{-1} \geq 1$  then,  $\left( \rho_k^{-1} - \sqrt{\frac{(1+2\rho_k^{-1})}{3}} \right) h(t) - (\rho_k^{-1} - 1) \leq 0$ .

Consequently,  $g(\cdot)$  is positive only if the second term of Equation 13 is negative, i.e.,  $\sqrt{\alpha \frac{\ln(t)}{u}} \leq$

$\frac{(\rho_k^{-1}-1)}{\left( \rho_k^{-1} + \sqrt{\frac{(1+2\rho_k^{-1})}{3}} \right)}$ . Since  $u \geq u_k$ , the last inequation is verified. Finally upper bounding Equation 11

for  $u \geq u_k$ :

$$\mathbb{P}(\{\phi_1(u_k)\}_{k,t}^\pi) \leq \sum_{u=u_k}^{t-1} e^{-\alpha \ln(u)/2} \leq \sum_{u=u_k}^{t-1} \frac{1}{u^{\alpha/2}} \leq \frac{1}{t^{\alpha/2-1}}$$

### D. Proof of Lemma 3

This proof follows the same steps as the the proof in Subsection VI-C.

From the definition of  $\{\phi_1(u_k)\}_{k,t}^\pi$  we can write that:  $\mathbb{P}(\{\phi_2\}_t^\pi) \leq \sum_{u=1}^{t-1} \mathbb{P}(B_{k^*,t}(u) \leq \mu^*)$

In the case of MUCB policies, we have:

$$\forall u \leq t, \mathbb{P}(B_{k^*,t}(u) \leq \mu^*) = \mathbb{P}\left(\bar{X}_{k^*,t}(u) \leq \frac{\mu^*}{M_{k^*,t}(u)}\right)$$

Consequently, we can upper bound the probability of occurrence of type 2 anomalies by:

$$\mathbb{P}(\{\phi_2\}_t^\pi) \leq \sum_{u=1}^{t-1} \mathbb{P}\left(\frac{\bar{X}_{k^*,t}(u)}{\mu^*} \leq \max\left\{0; \left(1 - \sqrt{\frac{\alpha \ln(t)}{T_{k,t}}}\right)\right\}\right)$$

Since  $\mu^* \max\left\{0; \left(1 - \sqrt{\frac{\alpha \ln(t)}{T_{k,t}}}\right)\right\} \leq \mu^*$  Cramer's condition is verified. Moreover since the machine is played when the maximal of the previous term is equal to 0, we can consider that  $u \geq \alpha \ln(t)$  and that:

$$\mu^* \max\left\{0; \left(1 - \sqrt{\frac{\alpha \ln(t)}{T_{k,t}}}\right)\right\} = \mu^* \left(1 - \sqrt{\frac{\alpha \ln(t)}{T_{k,t}}}\right)$$

Consequently, we can upper-bound the occurrence of Anomaly 2:

$$\mathbb{P}(\{\phi_2\}_t^\pi) \leq \sum_{u=\alpha \ln(t)}^{t-1} e^{-l_2(\beta_{k^*,t}(u))u} \quad (14)$$

Where,  $l_2(\beta_{k^*,t}(u))$  verifies the LDI as defined in Appendix VI-A. Thus, after mild simplifications we can write,

$$l_2(\beta_{k^*,t}(u)) \geq \frac{\frac{3\alpha \ln(t)}{u}}{2\left(1 + 2\left(1 - \sqrt{\frac{\alpha \ln(t)}{u}}\right)\right)} \geq \frac{\alpha \ln(t)}{2u}$$

Consequently, including this last inequality into Equation 14 ends the proof.